AF
ZFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

RAVI K. KAVURI et al.

Group Art Unit: 2163

Examiner: Cheryl M. Fernandes

Serial No.: 10/033,503

Filed: December 27, 2001

For: VIRTUAL VOLUME MANAGEMENT SYSTEM AND METHOD

Attorney Docket No.: 2001-028-NSC (STK 01028 PUS)

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an Appeal Brief for the appeal from the final rejection of claims 21-40 of the Office Action mailed November 28, 2005 for the above-identified patent application.

I. REAL PARTY IN INTEREST

The assignee of record in this application is Storage Technology Corporation ("Assignee"), a Delaware Corporation having a place of business at One StorageTek Drive, Louisville, CO, 80028, as set forth in the assignment recorded in the U.S. Patent and Trademark Office on July 23, 2002 at Reel 013107/Frame 0733.

CERTIFICATE OF MAILING UNDER 37 C.F.R. § 1.8 (FIRST CLASS MAIL)

I hereby certify that this paper, including all enclosures referred to herein, is being deposited with the United States Postal Service as first-class mail, postage pre-paid, in an envelope addressed to: Mail Stop Appeal Brief - Patents, Commissioner for Patents, U.S. Patent & Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450 on:

May 30, 2006	Jeffrey M. Szuma	Jeffrey M. Szuma
Date of Deposit	Name of Person Signing	Signature

Assignee Storage Technology Corporation is a wholly owned subsidiary of Sun Microsystems, Inc., a Delaware Corporation, having a place of business at 4120 Network Circle, Santa Clara, CA 95054.

II. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences known to the Appellants, the Appellants' legal representative, or the Assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 21-40 are pending in this application. Claims 1-20 have been canceled. Claims 21-40 have been rejected and are the subject of this appeal.

IV. STATUS OF AMENDMENTS

No amendment after the final rejection was filed.

V. SUMMARY OF CLAIMED SUBJECT MATTER

In order to improve data management, multiple physical storage devices, such as magnetic disk drives, can be grouped into sets or "pools." Pooling allows a set of specified physical disks to be abstracted as a single entity. Using pooling, when a logical or "virtual" disk is created, a user can specify a single pool of physical disks from which storage space is to be taken, rather than having to enumerate all physical disks that might be acceptable. In such a fashion, a single virtual disk may be presented to a user, while multiple pooled physical disks are specified and employed for actual storage of the user's data. (*See*, Specification, p. 5, ll. 11-27.)

It is known to use disk pooling in a Redundant Array of Independent/Inexpensive Disks (RAID). An exemplary disk pooling instance, such as may be

embodied in a RAID enclosure, is shown in Figure 1. A subset of physical disks (12A, 12B, 12C) is grouped or allocated to a pool (14). A virtual disk (16) is provided in communication with pool (14). As a result of the pooling, virtual disk (16) can obtain storage space only from physical disks (12A, 12B, 12C), and cannot obtain storage space from physical disk (12D). Figure 2 shows a RAID enclosure (20) that includes multiple physical disk drives (22i, 22ii, 22iii, . . . 22n), as well as an internal controller (24). Controller (24) pools the multiple disks (22i, 22ii, 22iii, . . . 22n) in order to present a single virtual disk (not shown). (*See*, Specification, p. 5, l. 28 - p. 6, l. 21.)

However, because RAID enclosures involve pooling of a fixed number of disks that are captive within the enclosure, storage capacity is only as extensible as the physical enclosure with its fixed number of disks allows. While larger RAID enclosures may be manufactured with more disks, a limit always exists on the number of disks that can ultimately be included. That is, an arbitrarily large physical enclosure is simply not possible. Similarly, while an existing RAID enclosure may be stocked with disks having greater storage capacity, it is not certain that a user's storage capacity requirements can continually be met by such "denser" RAID enclosures. (*See*, Specification, p. 6, l. 22 - p. 7, l. 2.)

The claimed invention is directed to automatic allocation of a storage device to a pool. More specifically, as an example, a user wishing to add a storage device in a storage area network need not identify which servers will use the device and then install the appropriate software driver on all such servers. Instead, the claimed invention automatically assigns the device to a pool, accounting for and handling any differences between storage devices so that such differences are not apparent to a user. As a result, a user need not be involved. (*See*, Specification; p. 7, l. 3 - p. 9, l. 13.)

In that regard, Figure 3 shows a storage area network (SAN) (32) that comprises a plurality of virtual storage volumes (34, 36) available to a user for use in storage and

retrieval of data. A controller (40) automatically allocates at least two of a plurality of network storage devices (*e.g.*, storage devices (38A, 38B, 38C) as shown in Figure 3) to a storage pool (42) and links at least one of the plurality of virtual storage volumes (*e.g.*, virtual storage volumes (34, 36) as shown in Figure 3) to the pool (42). Controller (40) performs such disk pooling inside of SAN (32) (*i.e.*, outside of a RAID enclosure or a server). As a result of the pooling, virtual storage volumes (34, 36) can obtain storage space only from network storage devices (38A, 38B, 38C), and cannot obtain storage space from network storage device (38D). In contrast to disk pooling in a RAID enclosure, however, whenever a user or an application requires additional disk storage space in the SAN, that need can be readily addressed by the addition of one or more additional storage devices to the appropriate pool. (*See*, Specification, p. 9, l. 14 - p. 10, l. 14.)

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 21-40 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,745,207 issued to Reuter et al. ("the '207 patent").

VII. ARGUMENT

The Rejection Of Claims 21-40 Under 35 U.S.C. § 102(e) In View Of The '207 Patent Should Be Reversed

The Appellants believe that claims 21-40 are not anticipated by the '207 patent. Anticipation is established only when a single prior art reference discloses each and every element of a claimed invention. *RCA Corp. v. Applied Digital Data Sys., Inc.*, 730 F.2d 1440, 1444, 221 USPQ 385, 388 (Fed. Cir. 1984). That is, there must be no difference between the claimed invention and the disclosure of the reference, as viewed by a person of ordinary skill in the field of the invention. *Scripps Clinic & Research Found. v. Genentech Inc.*, 927 F.2d 1565, 1576, 18 USPQ2d 1001, 1010 (Fed. Cir. 1991).

The Appellants believe that a *prima facie* case of anticipation has not been established for the rejection of claims 21-40. In that regard, independent claims 21 and 26 are directed to a system and method for managing a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data. As recited in claim 21, the system comprises a pool linked to at least one of the plurality of virtual storage volumes, and a controller for automatically allocating a storage device to the pool. As recited in claim 26, the method comprises automatically allocating a storage device to a pool and linking at least one of the plurality of virtual storage volumes to the pool.

Independent claims 31 and 36 are directed to a virtual volume management system and method. As recited in claim 31, the system comprises a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data, a storage pool linked to at least one of the plurality of virtual storage volumes, and a controller for automatically allocating a storage device to the pool. As recited in claim 36, the method comprises providing a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data, and automatically allocating a storage device to a storage pool and linking at least one of the plurality of virtual storage volumes to the pool.

Thus, as previously described in the above example, a user wishing to add a storage device in a storage area network need not identify which servers will use the device and then install the appropriate software driver on all such servers. Instead, the claimed invention automatically assigns the device to a pool, accounting for and handling any differences between storage devices so that such differences are not apparent to a user. As a result, a user need not be involved.

The '207 patent is directed to a system and method for managing virtual storage. However, those sections of the '207 patent cited by the Examiner disclose the creation of a virtual disk by allocating capacity in a storage pool to the new virtual disk. That is, existing

capacity in a storage pool is reserved for the newly created virtual disk. The storage pool from which existing capacity will be set aside for the new virtual disk may be selected manually or automatically. (*See, e.g.*, The '207 Patent; col. 12, ll. 18-38; col. 5, ll. 1-8.)

Thus, the '207 patent discloses reserving existing capacity in a storage pool for a newly created virtual disk. The '207 patent therefore fails to teach or suggest automatically allocating a storage device to a pool, as recited in various forms in independent claims 21, 26, 31 and 36. As a result, for at least the foregoing reasons, the Appellants believe that independent claims 21, 26, 31 and 36 are not anticipated by the '207 patent. There being no *prima facie* case of anticipation, the Appellants respectfully request that the final rejection of those claims under 35 U.S.C. §102(e) be reversed.

Claims 22-25, 27-30, 32-35 and 37-40 depend either directly or indirectly from independent claims 21, 26, 31 and 36, respectively, and include all the limitations thereof. As a result, and for at least the reasons set forth above concerning independent claims 21, 26, 31 and 36, the Appellants believe that claims 22-25, 27-30, 32-35 and 37-40 also are not anticipated by the '207 patent. Accordingly, the Appellants respectfully request that the final rejection of those claims under 35 U.S.C. §102(e) also be reversed.

CONCLUSION

In view of the foregoing, the Appellants respectfully request that the Board reverse the final rejection of claims 21-40 under 35 U.S.C. §102(e) as anticipated by the '207 patent.

The Commissioner is hereby authorized to charge the Appeal Brief fee of $500, as applicable under the provisions of 37 C.F.R. § 41.20(b)(2), and any other fee deficiency incurred as a result of the filing of this paper, to Deposit Account No. 19-4545. A duplicate copy of this page is enclosed for that purpose.

Respectfully submitted,

Ravi K. Kavuir, et al.

By: 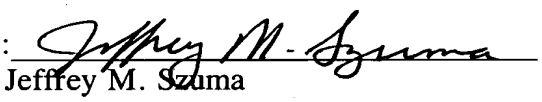

Jeffrey M. Szuma
Registration No. 35,700
Attorney for Applicant

Date: May 30, 2006

BROOKS KUSHMAN P.C.
1000 Town Center, 22nd Floor
Southfield, MI 48075-1238
Phone: 248-358-4400
Fax: 248-358-3351

Enclosure - Appendices

VIII. CLAIMS APPENDIX

1-20. (Canceled)

21. (Previously Presented) A system for managing a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data, the system comprising:

a pool linked to at least one of the plurality of virtual storage volumes; and

a controller for automatically allocating a storage device to the pool.

22. (Previously Presented) The system of claim 21 wherein the storage device comprises at least one physical disk.

23. (Previously Presented) The system of claim 21 wherein the storage device comprises at least one redundant array of independent disks.

24. (Previously Presented) The system of claim 23 wherein the storage device comprises at least one virtual storage volume.

25. (Previously Presented) The system of claim 24 wherein the at least one virtual storage volume comprises at least one virtual disk.

26. (Previously Presented) A method for managing a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data, the method comprising:

automatically allocating a storage device to a pool and linking at least one of the plurality of virtual storage volumes to the pool.

27. (Previously Presented) The method of claim 26 wherein the storage device

comprises at least one physical disk.

28. (Previously Presented) The method of claim 26 wherein the storage device comprises at least one redundant array of independent disks.

29. (Previously Presented) The method of claim 28 wherein the storage device comprises at least one virtual storage volume.

30. (Previously Presented) The method of claim 29 wherein the at least one virtual storage volume comprises at least one virtual disk.

31. (Previously Presented) A virtual volume management system comprising:

a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data;

a storage pool linked to at least one of the plurality of virtual storage volumes; and

a controller for automatically allocating a storage device to the pool.

32. (Previously Presented) The system of claim 31 wherein the storage device comprises at least one physical disk.

33. (Previously Presented) The system of claim 31 wherein the storage device comprises at least one redundant array of independent disks.

34. (Previously Presented) The system of claim 33 wherein the storage device comprises at least one virtual storage volume.

35. (Previously Presented) The system of claim 34 wherein the at least one

virtual storage volume comprises at least one virtual disk.

36. (Previously Presented) A virtual volume management method comprising:

providing a plurality of virtual storage volumes available to a user for use in storage and retrieval of user data; and

automatically allocating a storage device to a storage pool and linking at least one of the plurality of virtual storage volumes to the pool.

37. (Previously Presented) The method of claim 36 wherein the storage device comprises at least one physical disk.

38. (Previously Presented) The method of claim 36 wherein the storage device comprises at least one redundant array of independent disks.

39. (Previously Presented) The method of claim 38 wherein the storage device comprises at least one virtual storage volume.

40. (Previously Presented) The method of claim 39 wherein the at least one virtual storage volume comprises at least one virtual disk.

IX. EVIDENCE APPENDIX

None

X. RELATED PROCEEDINGS APPENDIX

None